



US 20010012295A1

(19) **United States**(12) **Patent Application Publication** (10) **Pub. No.: US 2001/0012295 A1**
WILFORD (43) **Pub. Date: Aug. 9, 2001**(54) **ENHANCED INTERNET PACKET ROUTING
LOOKUP****Publication Classification**(76) **Inventor: BRUCE A WILFORD, LOS ALTOS,
CA (US)**(51) **Int. Cl.⁷ H04L 12/28; H04L 12/56**(52) **U.S. Cl. 370/392; 370/390****Correspondence Address:****SWERNOFSKY LAW GROUP****P O BOX 390013****MOUNTAIN VIEW, CA 94039-0013 (US)**(57) **ABSTRACT**

The invention provides a method and system for routing information lookup for packets using routing protocols such as IP or IP multicast (IGMP), or MPLS (multiprotocol label switching, also known as tag switching or label switching) or MPLS multicast, in which both the destination address and the length of that destination address are matched using a lookup table having separate entries for selected addresses and their lengths, which can be accessed in parallel for multiple representations of destination addresses, and in which the input interface associated with the packet is included in the information used for matching.

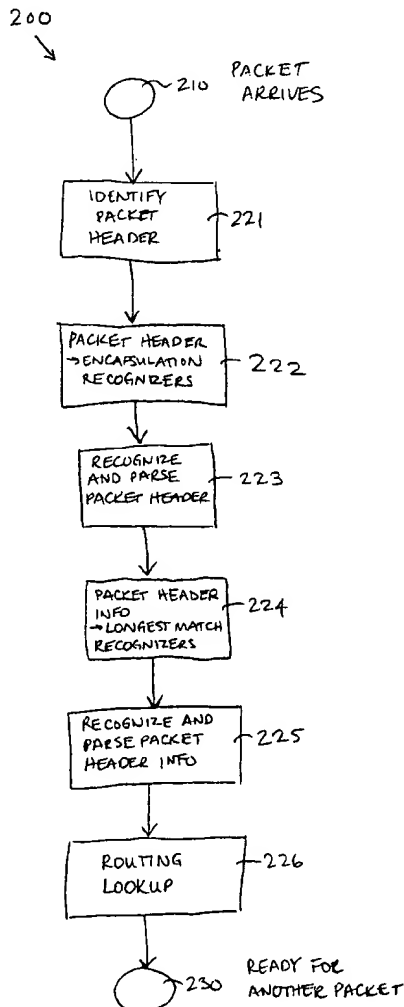
(*) **Notice:** This is a publication of a continued prosecution application (CPA) filed under 37 CFR 1.53(d).(21) **Appl. No.: 08/918,506**(22) **Filed: Aug. 22, 1997**

FIG. 1


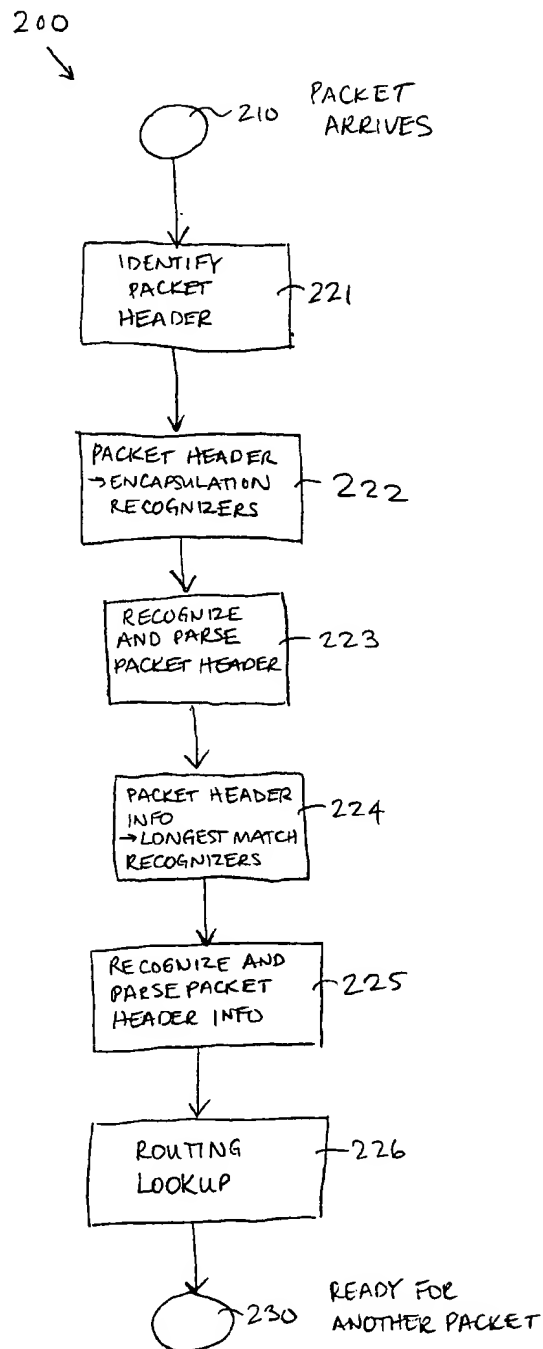


FIG. 2



ENHANCED INTERNET PACKET ROUTING LOOKUP

BACKGROUND OF THE INVENTION

[0001] 1. Field of the Invention

[0002] The invention relates to packet switching.

[0003] 2. Related Art

[0004] In a packet-switched network, a "router" is a device which receives packets on one or more input interfaces and which outputs those packets on one of a plurality of output interfaces, so as to move those packets within the network from a source device to a destination device. Each packet includes header information which indicates the destination device (and other information), and the router includes routing information which associates an output interface with information about the destination device (possibly with other information). The router can also perform other operations on packets, such as rewriting the packets according to their routing protocol or to reencapsulate the packets from a first routing protocol to a second routing protocol. It is advantageous for routers to operate as quickly as possible, so that as many packets as possible can be switched in a unit time.

[0005] One problem which has arisen in the art is that routing lookup, that is, determining an output interface on which to output the packet in response to the header information from the packet, can take substantial processing time, and that the demand on the router from that processing time continues to grow with increased network demand and increased network size. This problem has been exacerbated by addition of more complexity in routing protocols and routing techniques, such as the use of access control lists, network flow switching, quality of service requirements, and of attempts to determine the nature or routing treatment of packets in response to their source.

[0006] Some known routers, such as those described in U.S. application Ser. No. 08/655,429, "Network Flow Switching and Flow Data Export", filed May 28, 1996, in the name of inventors Darren Kerr and Barry Bruins, and assigned to Cisco Systems, Inc., attorney docket number CIS-016, and U.S. application Ser. No. 08/771,438, having the same title, filed Dec. 19, 1996, in the name of the same inventors, assigned to the same assignee, attorney docket number CIS-017, determine routing lookup in response to a network "flow" associated with the packet, rather than just in response to the destination address for the packet. This technique, called network flow switching, allows the router to determine routing procedure a single time for the entire flow, and to route all packets for the flow in the same manner.

[0007] While such routers are able to reduce the number of routing lookup operations required for a sequence of packets in a flow, they do not take advantage of all information available for rapid routing. For a first example, for some packets the amount of packet header information for determining an output interface can be relatively longer or shorter (as in IP protocols, where entire sets of destination addresses can sometimes be assigned to a single output interface). For a second example, for some packets (as in multicast proto-

cols, due to the nature of the multicast tree) the choice of output interface can be responsive to the input interface from which the packet was received.

[0008] Accordingly, it would be desirable to provide a method and system for performing routing lookup, which is responsive both to the header information associated with a packet and to the length of that header information, and which makes use of all useful information including an input interface from which the packet was received. These advantages are achieved in an embodiment of the invention in which both the destination address and the length of that destination address are matched using a lookup table having separate entries for selected addresses and their lengths, which can be accessed in parallel for multiple representations of destination addresses, and in which the input interface associated with the packet is included in the information used for matching.

SUMMARY OF INVENTION

[0009] The invention provides a method and system for routing information lookup for packets using routing protocols such as IP or IP multicast (IGMP), or MPLS (multi-protocol label switching, also known as tag switching or label switching) or MPLS multicast, in which both the destination address and the length of that destination address are matched using a lookup table having separate entries for selected addresses and their lengths, which can be accessed in parallel for multiple representations of destination addresses, and in which the input interface associated with the packet is included in the information used for matching.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 shows information for an example packet using an MPLS protocol.

[0011] FIG. 2 shows a process flow diagram of a method of operating the system.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0012] In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. Those skilled in the art would recognize after perusal of this application that embodiments of the invention can be implemented using general purpose processors or special purpose processors or other circuits adapted to particular process steps and data structures described herein, and that implementation of the process steps and data structures described herein would not require undue experimentation or further invention.

[0013] Inventions described herein can be used in conjunction with inventions described in the following applications:

[0014] application Ser. No. _____, filed the same day, Express Mail Mailing No. _____, in the name of the same inventor, titled "Multiprotocol Packet Recognition and Switching", attorney docket number CIS-020A; and

[0015] application Ser. No. _____, filed the same day, Express Mail Mailing No. _____, in the name of the same inventor, titled "Multiple Parallel Packet Routing Lookup", attorney docket number CIS-020B.

[0016] Each of these applications is hereby incorporated by reference as if fully set forth herein.

[0017] System For Enhanced Internet Packet Routing Lookup

[0018] A system for enhanced internet packet routing lookup includes a multiprotocol packet recognizer as further described in application Ser. No. _____, titled "Multiprotocol Packet Recognition and Switching", attorney docket number CIS-020A, and in application Ser. No. _____, titled "Multiple Parallel Packet Routing Lookup", attorney docket number CIS-020B, each of which is hereby incorporated by reference as if fully set forth herein. As more fully described in these incorporated applications, each packet is input to a packet buffer, coupled to a set of encapsulation recognizers and an encapsulation multiplexer, and coupled to a set of longest-match recognizers and a set of longest-match multiplexers, so as to first recognize the type of packet encapsulation and to second recognize relevant packet header information for routing the packet.

[0019] FIG. 1 shows information for an example packet using an MPLS protocol.

[0020] As used herein, "MPLS" refers to multiprotocol label switching, also known as "tag switching" or "label switching", as defined in RFC 2105, "Cisco Systems' Tag Switching Architecture Overview", by Y. Rekhter, D. Katz, E. Rosen, and G. Swallow (February 1997), which document describes label formats and packet formats, and is hereby incorporated by reference as if fully set forth herein.

[0021] 1. MPLS Packet Header

[0022] The information for the example packet includes an example packet header 100. In the case where the packet is encapsulated using an MPLS protocol, the packet header 100 includes a set of line information 101, a destination address (DA) 102, a source address (SA) 103, a fixed code 104 indicating that the packet is encapsulated using MPLS, and an MPLS label 105, followed by a sequence of words in continuation of the packet.

[0023] In a preferred embodiment, the line information 101 includes information about how the packet was received, including an input interface identifier 106 for an input interface associated with the packet.

[0024] In a preferred embodiment, the fixed code 104 indicating that the packet is encapsulated using MPLS comprises the value hexadecimal 88 47 (for MPLS unicast) or the value hexadecimal 88 48 (for MPLS multicast).

[0025] The MPLS label 105 includes information as specified by the MPLS protocol.

[0026] In a preferred embodiment, additional information is included with the packet header 100 when the packet header 100 is coupled to the longest-match recognizers, including a checksum value (CHK) 107 associated with the packet header 100 and computed by a checksum element.

[0027] Computation of the checksum value CHK 107 is further described in application Ser. No. _____, titled "Multiprotocol Packet Recognition and Switching", attorney docket number CIS-020A, and in application Ser. No. _____, titled "Multiple Parallel Packet Routing Lookup", attorney docket number CIS-020B, each of which is hereby incorporated by reference as if fully set forth herein. As more fully described in these incorporated applications, the checksum value CHK 107 includes one bit for each type of packet header protocol recognized by the encapsulation recognizers. In a preferred embodiment, this includes one bit for a correct IP checksum and correct packet header length, one bit for a correct CLNS checksum and correct packet header length, and one bit for just a correct packet header length. Since the MPLS protocol does not specify a checksum value, the latter bit is checked only for indicating correct packet header length.

[0028] 2. Packet Header Information

[0029] The information for the example packet includes a set of packet header information 110 gleaned by operation of the encapsulation recognizers and encapsulation multiplexer in response to the example packet header 100.

[0030] In a preferred embodiment, the packet header information 110 includes the checksum value CHK 107, the input interface identifier 106, a 20-bit label value 111, a three-bit class of service (COS) value 112, a one-bit top of stack (TOS) value 113, and a eight-bit time-to-live (TTL) value 114.

[0031] As described herein, the input interface identifier 106 comprises a value which identifies the input interface associated with the packet (that is, on which the packet was received).

[0032] As described herein, the checksum value CHK 107 is computed by the checksum element and simply confirms that a length of the packet header 100 in bytes is the same as specified in the packet header 100 itself.

[0033] As specified by the MPLS protocol, the 20-bit label value 111 includes information for routing the packet to an output interface.

[0034] As also specified by the MPLS protocol, other values comprising in the packet header information 110, including the class of service value COS 112, the type of service value TOS 113, and the time-to-live value TTL 114, are relevant to treatment of the packet but in a preferred embodiment do not effect the choice of output interface to which to route the packet. However, in alternative embodiments, the other values may be used to distinguish between different output interfaces which are coupled to differing routes to the destination. For example, the class of service value COS 112 might be used to distinguish between a relatively high-priority traffic route and a relatively low-priority traffic route.

[0035] 3. MPLS Routing Information

[0036] The information for the example packet includes a set of example routing information 120 gleaned by operation of the longest-match recognizers and longest-match multiplexers in response to the example packet header information 110.

[0037] The routing information 120 includes the input interface identifier 106 and the label value 111. For packets which use an MPLS multicast protocol, the input interface identifier 106 is needed for routing the packet; for packets which use an MPLS unicast protocol, the input interface identifier 106 is generally not needed for routing the packet.

[0038] Method of Operation

[0039] FIG. 2 shows a process flow diagram of a method of operating the system.

[0040] A method 200 of operating the system is controlled by elements described in the application Ser. No. _____, titled "Multiprotocol Packet Recognition and Switching", attorney docket number CIS-020A, and in application Ser. No. _____, titled "Multiple Parallel Packet Routing Lookup", attorney docket number CIS-020B, each of which is hereby incorporated by reference as if fully set forth herein. The operation of those elements is fully described therein. The method 200 includes flow points and process steps as described herein.

[0041] At a flow point 210, a packet using an MPLS protocol has arrived at an input interface, and is ready for processing.

[0042] At a step 221, a packet header 100 for the packet is identified and coupled to the packet buffer.

[0043] At a step 222, the packet header 100 is coupled from the packet buffer to the encapsulation recognizers. In the case where the packet is encapsulated using an MPLS protocol, the packet header 100 includes the packet header information 110 described herein with reference to FIG. 1. In the case where the packet is encapsulated using an MPLS protocol, the packet header 100 is recognized by at least one encapsulation recognizer configured to recognize those packets which are so encapsulated.

[0044] At a step 223, the at least one encapsulation recognizer configured to recognize packets using the MPLS protocol, along with the encapsulation multiplexer, collectively recognize and parse the packet header 100, so as to produce the packet header information 110.

[0045] At a step 224, the packet header information 110 is coupled from the encapsulation multiplexer to the longest-match recognizers. In the case where the packet is encapsulated using an MPLS protocol, the packet header information 110 includes the routing information 120 described

herein with reference to FIG. 1. In the case where the packet is encapsulated using an MPLS switching protocol, the packet header information 110 is recognized by at least one longest-match recognizer configured to recognize MPLS protocol information.

[0046] At a step 225, the at least one longest-match recognizer configured to recognize the MPLS protocol information, along with at least one longest-match multiplexer, collectively recognize and parse the packet header information 110, so as to produce the routing information 120.

[0047] At a step 226, the routing information 120 is used for routing lookup for the packet. The routing lookup is responsive to the routing information 120, and is thus responsive to both the input interface identifier 106 and to the label value 111.

[0048] At a flow point 230, routing of the packet has been accomplished, and the system is ready to route another packet.

[0049] Alternative Embodiments

[0050] Although preferred embodiments are disclosed herein, many variations are possible which remain within the concept, scope, and spirit of the invention, and these variations would become clear to those skilled in the art after perusal of this application.

1. A method, including the steps of receiving a packet header, said packet header including header information;

determining a packet treatment identifier responsive to said header information, said packet treatment identifier including an input interface identifier;

using said packet treatment identifier to access a memory having information regarding treatment of packets, said information being responsive to said input interface identifier.

2. A method as in claim 1, wherein said memory includes a plurality of entries matching said header information and differing with regard to said input interface identifier.

3. A method as in claim 2, wherein said memory includes a plurality of entries matching said header information and responsive to a multicast packet.

4. A method as in claim 3, wherein said multicast packet is encapsulated using an MPLS protocol.

* * * * *